NO SAFE HAVEN FOR THE VIRTUOUS

ABSTRACT:

In order to deal with the problem caused by environmental luck some proponents of robust virtue epistemology have attempted to argue that in virtue of satisfying the ability condition one will satisfy the safety condition. Call this idea the entailment thesis. In this paper it will be argued that the arguments that have been laid down for the entailment thesis entail a wrong kind of safety condition, one that we do not have in mind when we require our beliefs to be safe from error in order for them to be knowledge.

1. INTRODUCTION

The central thesis of robust virtue epistemology is that knowledge is a cognitive success that is attributable to the subject's cognitive abilities (Greco 2010; Sosa 2007; Zagzebski 1996). Robust virtue epistemologists aim to define knowledge solely in virtue-theoretic terms. Perhaps the most pressing argument against robust virtue epistemology is that it delivers wrong verdicts in cases featuring environmental luck. Accordingly, some epistemologists have abandoned robust virtue epistemology in favour of modest virtue epistemology which supplements the ability condition with an anti-luck condition that is designed to rule out all cases of knowledge undermining luck (Kelp 2013; Pritchard 2012). The safety condition is the anti-luck condition that modest virtue epistemologists are drawn to. According to the safety condition, a subject knows that *p* only if the subject could not easily have erred in her belief that *p*. The principal advantage that modest

1

virtue epistemology has over robust virtue epistemology is that it is able to deal with cases featuring environmental epistemic luck. The problem with modest virtue epistemology is that it is unnecessarily complex, at least compared to the extremely elegant robust virtue epistemology. Firstly, by adding an anti-luck condition to deal with cases of environmental luck the proponent of modest virtue epistemology owes us a story as to why knowledge has this structure: Why is it the case that the virtue-theoretic condition is able to deal with almost all of the cases but not quite (Pritchard 2010: 28)? Secondly, and more importantly, by retreating to modest virtue epistemology one will abandon the idea that knowledge is essentially a cognitive achievement that is creditable to the subject. Therefore modest virtue epistemologists can no longer explain the added value that knowledge has over mere true belief by referring to the idea that achievements are distinctly valuable, as has been argued by robust virtue epistemologists such as Zagzebski (2003), Greco (2010) and Sosa (2007). Pritchard (2010: 46) himself acknowledges this shortfall of modest virtue epistemology. If robust virtue epistemology could deal with cases featuring environmental luck, then it would clearly be preferable to modest virtue epistemology.

In response to the problem caused by environmental epistemic luck, some proponents of robust virtue epistemology have attempted to argue that in virtue of satisfying a properly understood ability condition one will satisfy the safety condition. Call this idea the *entailment thesis*. These robust virtue epistemologists think that the safety condition is a necessary condition for knowledge, but one that need not be referred to in the analysis of knowledge, since the ability condition already entails the safety condition.

In this paper it will be argued that the arguments that have been offered in support of the entailment thesis entail a wrong kind of safety condition, one that we do not have in mind when we require a belief to be safe from error in order for it to be knowledge. In the next section I will

lay out robust virtue epistemology and the problem of environmental luck. Two early solutions

to the problem are also canvassed. In the third section, four arguments for the entailment thesis

are presented and found wanting.

2. ROBUST VIRTUE EPISTEMOLOGY AND THE PROBLEM OF ENVIRONMENTAL

    LUCK

Robust virtue epistemologists maintain that knowledge is a cognitive achievement that is

creditable to the subject. More explicitly, they claim that the difference between knowledge and

ignorance resides in whether the subject's true belief is formed in such a way that it satisfies the

following *ability condition*:

A subject S's cognitive success is properly attributable to her cognitive abilities (that is, she

believes the truth because of her cognitive abilities, or her acquiring a true belief manifests her

cognitive abilities).[1]

Robust virtue epistemologists must claim that the truth of the subject's belief has to be

attributable to the subject's cognitive abilities, since otherwise they would not be able to deal

with Gettier cases. After all, in Gettier cases the subject gains a true belief and exhibits relevant

cognitive abilities, but nevertheless does not reach the truth in virtue of the fact that her belief

was competently formed. Rather, the subject reaches the truth in virtue of good epistemic luck

(Zagzebski 1996: 296-297). Armed with the ability condition the robust virtue epistemologist is

able to deal with Gettier cases. To see this, consider the following Gettier case:

---

[1] There are several ways to unpack the term 'attribution'. Sosa (2007, 2009) and Turri (2011) think that the truth of one's belief is attributable to one's cognitive abilities if it manifests one's cognitive abilities. Greco (2010, 2012), on the other hand, thinks that the truth of one's belief is attributable to one's cognitive abilities just in case the fact that one used one's cognitive abilities in acquiring the belief, is the best causal explanation for the fact why one gained a true, rather than a false belief. The former account is a metaphysically loaded account while the latter draws on explanatory considerations.

RODDY: Using his reliable perceptual faculties, Roddy noninferentially forms a true belief that there is a sheep in the field before him. His belief is also true. Unbeknownst to Roddy, however, the truth of his belief is completely unconnected to the manner in which he acquired this belief since the object he is looking at in the field is not a sheep at all, but rather a sheep-shaped object which is obscuring from view the real sheep hidden behind (Pritchard 2012: 251).

Even though Roddy's belief is true and the product of his cognitive abilities, his cognitive success is not attributable to his cognitive abilities. After all, Roddy didn't even see the sheep. Hence he does not know that there is a sheep in the fields.

In RODDY the epistemic luck at play is intervening luck. Something quite literally gets 'in between' the subject's cognitive abilities and the truth of his belief. Since the ability condition is not satisfied in standard Gettier cases, such as RODDY, robust virtue epistemologists can deal with such cases. The problem emerges, however, when we consider cases that feature environmental epistemic luck, the classic example being the barn façade case:

BARNEY: Using his reliable perceptual faculties, Barney noninferentially forms a true belief that the object in front of him is a barn. Barney is indeed looking at a barn. Unbeknownst to Barney, however, he is in an epistemically unfriendly environment when it comes to making observations of this sort, since most objects that look like barns in these parts are in fact barn façades (Pritchard 2012: 251).

The mainstream intuition regarding this case is that it is a case of ignorance. After all, Barney could very easily have formed a false belief. Notice that the BARNEY case differs in important respects from normal Gettier cases. Barney actually sees the barn. Because Barney actually sees

the barn, the luck involved is not intervening, as in the RODDY case, but environmental. In cases of environmental luck the subject is hooked up with the reality in the right kind of way, but the environment in which the subject finds herself in ensures that the fact that the subject gained a true belief is a matter of luck (Pritchard 2010: 36). Pritchard has argued that the satisfaction of the ability condition is compatible with environmental luck. He writes that:

> Barney is, after all, really seeing a genuine barn, unlike, say, Roddy, who merely thinks that he is seeing a genuine sheep. In a very real sense, then, Barney's cognitive abilities are putting him in touch with the relevant fact unlike in standard Gettier-style cases, where there is a kind of fissure between ability and fact, albeit one that does not prevent the agent from having a true belief regardless […]. The problem, however, is that given that Barney does undertake, using his cognitive abilities, a genuine perception of the barn, it seems that his cognitive success is explained by his cognitive abilities, unlike in standard Gettier-style cases (2012: 267).

If Pritchard is right about this, then robust virtue epistemology faces a serious challenge: it is unable to explain why subjects in cases featuring environmental luck lack knowledge.

Robust virtue epistemologists have tried to deal with cases featuring environmental luck in two ways. Sosa has acknowledged that the truth of one's belief can be attributable to one's cognitive agency in cases of environmental luck. Sosa produces an error theory according to which the subject gains 'animal knowledge' but lacks 'reflective knowledge' in cases featuring

environmental luck (2007: 35-36). The problem with this response is that Barney does not seem to gain even animal knowledge.[2]

Greco, unlike Sosa, is not ready to bite the bullet regarding the BARNEY case. Instead he argues that abilities are defined in relation to environments, and that Barney lacks the relevant ability in the barn façade county since he does not have a high rate of success in acquiring true beliefs about barns in the environment in which he is across the scope of nearby possible worlds (2010: 76-77). Therefore Barney cannot believe the truth because his belief is the product of his cognitive abilities, since in fact Barney does not have the relevant ability.

Even though this response superficially seems to be able to deal with the BARNEY case, it rests on a wrongheaded understanding of what cognitive abilities are. It is plausible that abilities should in some sense be relativized to environments, but it seems odd to say that simply by being in a certain environment one cannot use some specific ability, even if the environment does not interact or hinder one's actions in any way. Pritchard (2010: 38) has offered a case that illustrates this nicely: if I have the ability to play the piano, but not the ability to play the piano underwater, and I happen to be in a room that could at any moment be engulfed by water, then am I practising a different ability than I would be if the room were not about to be engulfed by water soon? Intuitively, I am practising the very same piano playing ability that I use in more ordinary circumstances. Greco's definition of what abilities are gets Pritchard's pianist case backwards. Moreover, the reason why it gets it backwards is the same reason why it is able to deal with the

---

[2] Greco (2016: 59-60) highlights the problems of thinking that one could gain animal knowledge in such a case. In his later work, Sosa has emphasized that our intuitions regarding knowledge often track reflective knowledge rather than mere animal knowledge. Crucially, Sosa nowadays holds that reflective knowledge, or "knowledge full well", entails a formulation of the safety condition. We will return to Sosa's present view in the final section.

BARNEY case. Therefore the solution that Greco offers is not sufficiently general, and might be accused of being *ad hoc*.

But given that the outcome that Greco is striving for (i.e. trying to deliver the verdict of ignorance regarding cases like BARNEY) is clearly preferable to the outcome that Sosa promises, it is not surprising that many robust virtue epistemologists have come to Greco's aid. Interestingly, many of these authors argue that by satisfying a properly formulated ability condition, one will satisfy some sort of safety condition.

To introduce some terminology, following Littlejohn (2014: 373), let us label those who think that cognitive achievement is compatible with environmental luck *compatibilists*, and those who think that cognitive achievement is not compatible with environmental luck *incompatibilists*. The incompatibilists claim that by satisfying the virtue-theoretic condition one will satisfy a properly formulated anti-luck condition, namely some formulation of the safety principle, which will rule out all species of harmful epistemic luck. The compatibilists, on the other hand, maintain that either the virtue epistemologist must add a separate anti-luck condition to deal with cases of environmental luck, or they should be content with the consequence that there is 'knowledge of a sort' present in cases that feature environmental luck. Since the safety condition is at the centre of the incompatibilist's argument, it will be useful to explicate the vague notion that we started with. Here is a mainstream formulation of the safety condition:

A subject S's true belief that *p* is *safe* if and only if,

(i)   in nearly all nearby possible worlds, and in all of the very closest nearby possible worlds in which S believes that *p*, (via the same method M that she uses in the actual world), *p* is true.

The safety condition is a promising candidate for being a necessary condition for knowledge.[3] It is an anti-luck condition which draws direct support from Pritchard's (2005) modal analysis of luck.[4] The safety condition is able to deal with standard Gettier cases such as RODDY, but more importantly for the purposes of the incompatibilist, the safety condition also gives the correct verdict in the BARNEY case. There are plenty of nearby possible worlds where Barney believes that the structure in front of him is a barn, while his belief is false, since he could be facing a mere façade. Thus, his belief is not safe.

Since the safety condition seems able to deal with all cases of knowledge undermining luck it is easy to see why the incompatibilist is drawn to the idea that by satisfying the ability condition one will satisfy the safety condition. If that entailment thesis were true, then robust virtue epistemologists would no longer be threatened by counterexamples that relied on environmental luck.

But if the safety condition is such a great anti-luck condition, one might wonder why not define knowledge simply as a safe belief? The reason why we should be cautious of making such an inference is that the safety condition is probably not a sufficient condition for knowledge. The safety condition is trivially satisfied in cases where the subject forms a belief in a necessary

---

[3] The safety condition has been advocated by many epistemologists as a necessary condition for knowledge. For different formulations of the condition, see Sainsbury (1997), Sosa (1999), Williamson (2000), Luper (2003), and Pritchard (2012). For recent critiques of the safety condition, see Neta and Rohrbaugh (2004), Comesaña (2005), Hawthorne and Lasonen-Aarnio (2009), Bogardus (2014), and Miracchi (2015).
[4] For critiques of Pritchard's analysis of luck see Coffman (2007), Riggs (2007), and Lackey (2008).

truth, since there are no possible worlds where the belief that the subject formed is false. But clearly all beliefs that are necessarily true do not amount to knowledge. If I happen to form a belief in a necessary truth by flipping a coin, then clearly I do not gain knowledge, even though my belief could not have been false.

The reason why the safety condition is susceptible to cases featuring necessary truths is because it only demands that the belief that the subject *actually* formed has to be true in the nearby possible worlds where the subject continues to hold that very same belief. But in order to truly be safe from error it is not enough that the belief that the subject formed in the actual world could not easily have been false. Rather, the subject must be safe from error as an epistemic agent, i.e. she must be in an epistemic position where she could not easily have gained a false belief in her inquiry. Therefore, proponents of the safety condition should globalize the notion of safety to a set of relevant propositions.[5] A globalized version of the safety condition is able to deal with cases featuring necessary truths. If I happen to form a true belief in a necessary truth by flipping a coin, I could easily have ended up having a false belief (albeit a different one that I ended up with in the actual world). Therefore I was not properly safe from error in forming my belief.

But a proponent of robust virtue epistemology need not make these alterations to the safety condition, because the truth of a subject's belief will not be attributable to the subject's cognitive abilities if she forms her belief on a deviant basis (e.g. by flipping a coin). Therefore, the virtue

---

[5] Globalized versions of the safety condition have been put forward by Williamson (2009), Pritchard (2012), Ball (2016) and Hirvelä (2017a, 2017b).

theoretic condition and the safety condition supplement each other in a harmonious way.[6] If the incompatibilists are right, then we have a very powerful and simple theory of knowledge at hand which contains only one necessary and sufficient condition, unlike the more cumbersome anti-luck virtue epistemology advocated by Pritchard. Now we will turn to look at the arguments for the incompatibilist position.

3. FROM VIRTUE TO SAFETY

The idea that genuine achievement rules out success by luck seems compelling. After all, if one's success is to count as an achievement, then the fact that one succeeds has to be attributable to one's relevant abilities to a sufficient degree. If one is inclined to understand the attribution relation in the way that Greco does, then we will have a straightforward answer as to why achievements are incompatible with luck. According to Greco, in cases of knowledge one's cognitive abilities have to be the best causal explanation for the fact that one believes the truth (2010: 12, 74). If one's cognitive abilities are the best causal explanation for the fact that one believes the truth, then it seems that one's belief cannot be true simply as a matter of luck, since if it were, then one's cognitive abilities would not be the best explanation as to why one gained a true, rather than a false belief.

The problem that arises from demanding that one's cognitive abilities have to be the best causal explanation for the fact that one gained a true belief is that it places too strong conditions on knowledge. More precisely, by taking the line of defence that Greco takes, the virtue epistemologist will be ill-suited to deal with cases of testimonial knowledge, where intuitively

---

[6] In fact this is the point of anti-luck virtue epistemology which is advocated by Pritchard (2012). Pritchard however thinks that the ability and anti-luck conditions are two separate conditions that need to be combined, and thus he is a compatibilist.

the best causal explanation for the subject's true belief is not the subject's cognitive abilities, but the testifier's abilities. Lackey (2007, 2009) has argued that this forces the virtue epistemologist into a dilemma; (i) either the virtue epistemologist can strengthen the ability condition and allow herself to be able to deal with cases featuring environmental luck while handicapping herself in the face of cases featuring testimonial knowledge, (ii) or the virtue epistemologist can weaken the ability condition and be able to deal with cases of testimonial knowledge, but as a result be unable to deal with Gettier cases. If these really are the available options, then the prospects of robust virtue epistemology are dim indeed.

But perhaps there is some other way to pursue the incompatibilist line of argument that does not get tangled in the dilemma proposed by Lackey. One natural way to proceed would be to note that luck and achievement come in degrees. The luckier a success is, the less creditable it is to the agent who accomplished it. Carter (2016) argues that by abandoning the rigid idea that the attribution of success to ability is incompatible with the attribution of that success to luck, we can salvage the virtue-theoretic programme. Since luck comes in degrees, so do the harmful effects that it has on knowledge. Accordingly, Carter proposes the following condition on when a success can be attributable to a subject's cognitive abilities:

> BALANCE: If the correctness of *S*'s belief that *p* depends (sufficiently) on *S*'s cognitive ability, then it depends on *S*'s ability more so than luck that *S*'s belief that *p* is true (2016: 147).

BALANCE respects the fact that safety (and luck) come in degrees. Therefore, it would seem that by adopting BALANCE the incompatibilist is at least in a position to argue for the entailment thesis in a way that avoids Lackey's dilemma. The next premises in Carter's argument for the entailment thesis are:

(i)   The correctness of S's belief depends more so on S's ability than luck, only if the worlds where S could not easily have erred (by employing the same cognitive abilities that she does in the actual world) are closer to the actual world than the worlds where S errs (while using the same cognitive abilities that she uses in the actual world).

(ii)   If the worlds in which S does not have a false belief are closer to the actual world than the worlds where she has a false belief (holding the abilities employed in the actual world fixed), then S's belief is safe in the appropriate sense.

From BALANCE, (i) and (ii) Carter derives the entailment thesis (2016: 148). Though it seems trivial that Carter can establish the entailment thesis from his premises, it is open to question what kind of safety condition the subject will satisfy by satisfying the ability condition in the way that Carter understands the ability condition.

The crucial premise in Carter's argument for the entailment thesis is (ii). Carter claims that if the worlds in which S is not led to a false belief by using the same cognitive abilities that she uses in the actual world are closer to the actual world than the worlds where she is led to a false belief by using those very same cognitive abilities, then S's belief is safe in the appropriate sense. This certainly is something that is usually the case when a subject's belief is safe. However, this

picture of what it takes for a belief to be safe is too lenient. To see this, consider a slightly altered version of a case given by Sosa (1999: 145).

RUBBISH CHUTE: S is about to drop a rubbish bag down the chute of her high-rise condo. Unfortunately, and unbeknownst to S, the chute has been damaged and the bag will end up in the basement only if S drops the bag in the way she usually does it. Therefore the bag will end up in the basement only in the scope of the very nearest possible worlds, but will get snagged on some imperfection in most nearby possible worlds. As S drops the bag she forms the true belief that the bag will soon be in the basement.

S's belief is clearly unsafe in this case. In most nearby possible worlds, where S continues to believe via the same method of belief-formation that she uses in the actual world that the bag will soon be in the basement, her belief will end up being false. Therefore she does not satisfy the safety condition. Moreover, it seems that the safety condition delivers the correct verdict regarding this case. The subject does not know that the bag will soon be in the basement. However, it is easy to see that the formulation of safety that is at play in Carter's argument for the entailment thesis is satisfied by the subject of the RUBBISH CHUTE case. After all, the worlds in which S forms a true belief by using the same cognitive abilities that she uses in the actual world are closer to the actual world than the worlds where she forms a false belief by using the very same cognitive abilities. Therefore S's belief is appropriately safe according to Carter. But it is quite clear that S's belief is not safe in the appropriate sense, and therefore the formulation of safety that Carter has in mind is not the one that we have in mind when we

require a belief to be safe from error in order for that belief to amount to knowledge. Carter's argument for the entailment thesis fails.

Another way to argue for the entailment thesis would be to claim that the subjects in cases featuring environmental luck have not been afforded with an opportunity to exhibit their cognitive abilities and the manifestation of their abilities requires that they be given an opportunity to practise their relevant abilities (Littlejohn 2014: 373). In the BARNEY case, Barney is deprived of the opportunity to exhibit his discriminatory abilities by ensuring that there is no distinct look on the basis of which to discriminate real barns from fakes. Littlejohn says:

> When there isn't such a basis, there is no basis for attributing successful classification to the subject's abilities… the subject might occasionally classify the odd F as an F on the basis of how it looks, but it doesn't seem that we could attribute success to something that the subject was sensitive to for the simple reason that there was nothing that the subject could have been sensitive to in responding to the way that an F looked that could have been the distinctive mark of an F (2014: 377).

Gaultier (2014: 481) argues along similar lines in favour of robust virtue epistemology. According to him Barney uses the same cognitive abilities in both the case that there are numerous fakes around him, and in the more mundane case, where there are no fakes around. But even though the subject gains the same true belief in both cases, it does not mean that the subject attains the same cognitive success in both cases.

14

According to Gaultier, this is because "strictly speaking, *there are no cognitive abilities to form true beliefs* – perceptual or otherwise. Perceptual abilities, for instance, are not abilities to *form true perceptual beliefs*, but abilities *to* (visually, auditorily, tactually, etc.) *identify, discriminate, detect or recognize such and such a thing in one's environment*" (2014: 482). In the BARNEY case, Barney does not succeed in identifying a barn. According to Gaultier, identifying, discriminating, etc. never happens in isolation: "to identify, is necessarily to identify within a perceptual context or environment" (Gaultier 2014: 482).[7]

If Gaultier and Littlejohn are right on this matter (and I suspect that they are), is the incompatibilist in a position to argue for the entailment thesis? Both Littlejohn and Gaultier think that once the true nature of cognitive abilities is unveiled, cases of environmental luck will no longer pose a threat to robust virtue epistemology. In fact, they both seem to think that the entailment thesis is true. Gaultier claims that "[f]or the truth of a perceptual belief to be credited to the believer, this belief has to be safe" (2014: 486), while Littlejohn maintains that "it looks like one must satisfy some sort of safety principle by virtue of satisfying the ability condition" (2014: 379). If we accept the reasoning of Littlejohn and Gaultier, then it seems that we are in a position to accept the entailment thesis. The famous BARNEY case is dismantled. It would seem that there are no cases where the truth of a belief is attributable to one's cognitive abilities and yet the belief is unsafe.

---

[7] Both Gaultier and Littlejohn draw a distinction between practical and doxastic abilities. For example, according to Littlejohn cognitive abilities "put one in touch with the things in one's surroundings" (2014: 375), while practical abilities do not. Because of the divergence between doxastic and practical abilities, Littlejohn and Gaultier can explain why cases of environmental luck are not analogous to the kind of archery cases that modest virtue epistemologists often cite.

However, as with Carter's solution, it is fair to ask "what kind of safety principle will one satisfy in virtue of satisfying the ability condition that Gaultier and Littlejohn have in mind?" Regrettably, the kind of safety condition that is entailed by satisfying the ability condition, as Gaultier and Littlejohn conceive it, is not the one we have in mind when we require a belief to be safe from error in order for it to be knowledge. To see this, consider the following case of environmental luck:

BOWL OF 'FRUIT': Sofia is looking at what seems to be a bowl of fruit and forms the true belief that "that is an apple" while pointing at an apple. Sofia has excellent eyesight and the perceptual conditions are fine. However, as it happens, only the apples in the bowl are real. The rest of the 'fruits' in the bowl are very realistic plastic replicas that would have fooled anyone. Moreover, due to various reasons, there could not have been any fake apples in the bowl. Luckily, Sofia happened to point at an apple.

Granted that perceptual abilities are abilities to discriminate, we can then ask: "Is there any distinct mark that can be used to identify apples from non-apples in this case?" If there is, then Sofia has been given an opportunity to exhibit her discriminatory abilities, as Littlejohn and Gaultier demand. But of course one of the relevant differences between the BOWL OF 'FRUIT' case and the BARNEY case is that there is such a feature. The apples have a distinguishing mark, namely they look like apples. There are no objects in the bowl that share this look that are not apples, nor could there be any, since it is stipulated in the case description that there could not have been any fake apples in the bowl. Therefore, the truth of Sofia's belief seems to be attributable to her cognitive abilities. Nevertheless the BOWL OF 'FRUIT' is just as clearly a

case of ignorance as the BARNEY case was. The kind of safety principle that might be entailed by the satisfaction of the ability condition as Gaultier and Littlejohn conceive it is not the one that we have in mind when we require a belief to be safe from error in order for it to be knowledge.

Moreover, it cannot be argued on behalf of Gaultier and Littlejohn that Sofia does not succeed in identifying an apple in the BOWL OF 'FRUIT' on the grounds that the perceptual environment in which she is in ensures that she is not able to distinguish fruits from non-fruits, since Sofia would have misclassified the replicas as fruit. The reason for this is that the fact that Sofia is in such a situation cannot preclude the possibility that she successfully identifies an apple. After all, it is possible to identify X as an apple without being able to identify it as a fruit. In the simplest case, one might lack the concept of fruit while possessing that of an apple. The fact that one lacks the concept of fruit should not in itself entail that one cannot recognize something as an apple. In a more interesting case, one possesses the relevant concepts, but does not know that one belongs to the other's extension. For instance, one might recognize that Y is a dolphin without recognizing it as a mammal. In an even more complex, case a subject might recognize Z as a pikeperch without being able to distinguish fish from non-fish in her environment since she would have misclassified the nearby dolphins as fish and the nearby stone fish as a stone. Since it is clearly possible to identify something as an apple without being able to distinguish fruit from non-fruit in one's environment, it cannot be argued that Sofia fails to identify an apple in the BOWL OF 'FRUIT' on the grounds that she is not able to discriminate fruit from non-fruit in her predicament.

I think that the mistake that Gaultier and Littlejohn made is understandable, since the mainstream formulation of the safety condition is also satisfied in the BOWL OF 'FRUIT' case. After all, in all the nearby possible worlds where Sofia continues to believe what she believes in the actual world her belief continues to be true, since there are no fake apples in the bowl. Sofia could not easily have falsely believed that [that's an apple]. In this respect the BOWL OF 'FRUIT' is analogous to cases featuring necessary truths.

Earlier we noted that the safety condition could easily be fixed to deal with cases featuring necessary truths. The solution was to globalize it to a set of propositions in which the subject could easily have ended up believing in her inquiry. Sofia could clearly have ended up with a false belief in her situation. She could easily have falsely believed that [that's a persimmon], [that's a banana], etc. Therefore she does not satisfy the safety condition once we understand it properly.

We also noted that proponents of robust virtue epistemology do not need to reformulate the safety condition in this manner, because the virtue-theoretic component of their view already allowed them to deal with cases featuring necessary truths. The BOWL OF 'FRUIT' case shows, however, that this latter point is false. Virtue epistemological theories cast along the lines that Gaultier and Littlejohn propose are unable to deal with cases like the BOWL OF 'FRUIT' precisely because the conception of safety that they rely on is inapt.

Finally, let us return to Sosa's present position.[8] In his most recent work, Sosa has returned to the idea that knowledge requires both virtue and safety, and that the satisfaction of the virtue-theoretic condition entails the satisfaction of the safety condition. More precisely, Sosa holds that "knowledge full well" entails safety. The reason why knowledge full well entails safety is because knowledge full well requires aptness of *judgment*. In order for a judgment that *p* to be fully apt, the subject must have situated herself correctly in her cognitive environment so that she knows that she would (likely enough) affirm correctly in her endeavour to judge correctly whether *p* (Sosa 2015: 79). Concerning the BARNEY case Sosa writes:

> Again, in order to know full well, Barney must know that if in his conditions he affirmed that he faces a barn, not easily would he be wrong. He needs knowledge of this conditional in order to guide himself to apt affirmation in the way required for *full* aptness of affirmation, which is what apt judgment requires. So, his judgment can then be apt only if safe. Accordingly, Barney knows full well only if his constitutive judgment is safe. (2015: 79)

For the sake of the argument, let us grant that knowledge full well entails safety, and that our intuitions regarding knowledge generally track knowledge full well rather than mere animal knowledge. The question we need to ask is what kind of safety condition is entailed by knowledge full well.

---

[8] I would like to thank an anonymous referee of *Dialectica* for urging me to examine Sosa's present position in greater detail.

It is important to note that the kind of safety condition that Sosa endorses has to be motivated by virtue epistemological considerations alone. Otherwise, his theory would not be a version of robust virtue epistemology. Thus we need to examine Sosa's view of what competences are in order to unearth the kind of safety condition he thinks is necessary for knowledge.

According to Sosa, competences are dispositions to succeed. Dispositions (and competences in particular) have a three-part structure: a seat, shape and situation. The seat of a visual competence is rods and cones, the shape is being awake and clearheaded, while the situation includes conditions such as being in adequate lighting conditions. A subject has complete competence only if she has the seat of the competence and is in proper shape and situation for the operation of that competence. A subject who is in possession of complete competence to ϕ would succeed (likely enough) in ϕ-ing if she attempted to ϕ while in possession of the complete competence. (Sosa 2015: 95-96)

Given that knowledge full well is defined with the help of complete competence, it is easy to see that if a subject is in possession of a complete epistemic competence with respect to a certain question, the subject would most likely believe the truth with respect to that question (if she formed a belief at all), while she retains her full competence. Therefore, Sosa relativizes the safety condition to the proper seat/shape/situation. Moreover, he recognizes that the subject must be safe not just with respect to the belief that she formed in the actual world, but that she has to be safe in her answer to a relevant question (2015: 52-53, 123). Therefore, he advocates a globalized version of the safety condition. Finally, Sosa clearly endorses a *weak* safety condition,

according to which a subject's belief is safe just in case in most nearby possible worlds where she believes in a similar enough belief her belief is true (2015: 123).

From these observations we can extract the safety condition that Sosa currently endorses:

SAFETY^S: S's true belief that *p* is safe just in case in most nearby possible worlds where S forms a similar enough belief, while she remains in possession of complete competence, her belief is true.[9]

SAFETY^S is a very promising safety condition. It is able to deal with the BOWL OF 'FRUIT' case because it is globalized to a set of propositions regarding which the subject has to be safe from error. Moreover, SAFETY^S is not satisfied in the RUBBISH CHUTE case since in most nearby possible worlds where the subject forms the belief that the rubbish bag is in the basement her belief is false. Finally, by endorsing SAFETY^S as a necessary condition for knowledge full well, Sosa is able to deliver the intuitive verdict with respect to the BARNEY case, since Barney could easily have formed a false belief in his epistemic circumstances.

Even though Sosa is able to dismantle all of these counterexamples with the help of SAFETY^S there are good reasons to think that SAFETY^S is not the kind of safety condition that we have in mind when we require that a subject must be safe from error in order to have knowledge. The reason for this is that SAFETY^S delivers absurd verdicts regarding lottery propositions or else breaks plausible epistemic principles.

---

[9] Greco (2016: 54) extracts a similar safety condition from Sosa's virtue epistemology, though he does not make it explicit that Sosa favours a globalized version of the safety condition.

A subject who believes that her lottery ticket has lost simply on the basis of the odds involved will satisfy Sosa's conditions for knowledge. After all, if I believed that my ticket has lost solely on the basis of the odds involved, my belief would continue to be true in most nearby possible worlds where held on the same basis. Sosa embraces this counterintuitive consequence and seeks to explain away the intuition that we cannot know lottery propositions.[10] Sosa produces two arguments as to why we in reality can know that our lottery ticket is a loser solely on the basis of the odds involved. Firstly, he introduces a thought experiment which he takes to be analogous to the lottery case where intuitively the subject acquires knowledge.

BAD APPLE:

> Apple has grown so large that a million operators stand by constantly to provide answers to any simple Apple question, and suppose technology randomizes in such a way that there are a million-minus-one operators no farther modally from me than the one I actually reach when I place my call. And suppose there is one bad apple in the lot, a liar who would give me an incorrect answer. But suppose further that Apple's service continues to be great, so that all the other million-minus-one operators are infallibly reliable in their answers to such questions. Is that enough to preclude me from learning the answer to my question from the operator I actually reach? (Sosa 2015: 119)

First of all, this case differs significantly from the lottery case. A crucial feature of the lottery case is that the subject forms the belief that her ticket is a loser solely on the basis of the odds involved. However, it is quite natural to think that in BAD APPLE the subject does not form her belief through probabilistic reasoning but via trusting the testimony of someone. At the very least

---

[10] The vast majority of epistemologists think that one cannot know that one's ticket has lost simply on the basis of the odds involved. There is far more disagreement about whether one can be justified in believing that a certain ticket has lost simply on the basis of the odds involved.

this creates "noise" that might muddle our intuitions regarding the case. Given how the case is described it is, for example, an open possibility that the subject would not trust the word of the bad apple, due to some feature of her testimony, and hence her belief could satisfy the safety condition as introduced in section 2. Even if we assume that the subject would trust the bad apple, and that she hence fails to satisfy the safety condition, there are features of the case that allow the proponent of the safety condition, as it was introduced in section 2, to offer a plausible story as to why we might mistakenly think that the subject knows. For example, given that the subject trusts the testimony of someone whom she has good reason to consider reliable and honest, her belief has properties which indicate that her belief amounts to knowledge, and hence one could argue that her belief is justified. This is not true of the lottery case, however, since purely probabilistic grounds can never indicate that a belief would satisfy the safety condition, nor could such a belief be potential knowledge.[11]

Finally, to make the case analogous to the lottery case, it would need to be specified that the sole reason why the subject formed a belief in response to the testimony she received from the apple operator is the probabilistic reasoning she undertook. Crucially, if it is made explicit in the case description that the subject believes what the operator tells her purely on probabilistic grounds, the intuition that the subject knows vanishes.

---

[11] Bird (2007) and Ichikawa (2014) have argued that justification is potential knowledge. While their views have subtle differences both agree that a belief in a lottery proposition can never be justified because it could never amount to knowledge. A belief based on testimony could clearly, however, amount to knowledge and hence can be justified.

The second reason that Sosa produces for his claim that we erroneously think that one cannot know that one's ticket is a loser solely on the basis of the odds involved is that we mistakenly think that knowledge requires sensitivity rather than safety (2015: 120).[12] Sosa writes:

> Sensitivity is noxious in its power to mislead. It misleads us not only with regard to skeptical scenarios, but also with regard to lottery scenarios. Sensitivity is easily confused with safety, because of how plausible it is that subjunctive conditionals contrapose. Once we are clear that they do not contrapose, we can then countenance basis-relative safety while discarding sensitivity. (2015: 120)

Even if Sosa is right that our intuitions regarding knowledge often mistakenly track sensitivity rather than safety, I think he has not fully appreciated the force of the lottery puzzle. The reason why epistemologists are hesitant to say that one could know that one's ticket has lost solely on the basis of the odds involved has not so much to do with intuitions, but rather with certain principles that knowledge ought to respect.

For the sake of argument, let us assume that I can know that my ticket has lost simply on the basis of the odds involved. Since there is nothing special about my ticket, I ought to be able to know that your ticket has lost as well. To say otherwise would make knowledge dependent on epistemically arbitrary factors. If I am able to know of our tickets that they have lost, then, by parity of reasoning I should be able to know of every single ticket that is a loser that is has lost.

---

[12] According to the sensitivity condition S's belief that *p* is sensitive just in case if *p* were false S would not believe that *p*. The sensitivity condition was originally proposed by Dretske (1971) and Nozick (1981) as a necessary condition for knowledge. A subject who believes that her lottery ticket has lost simply on the basis of the odds involved does not satisfy the sensitivity condition, because in the nearest possible world where she wins the lottery she will still continue to believe that she lost if she forms her belief on the basis of the odds involved.

That is, I should be able to know of every single ticket that lost that it has lost, but that is absurd! Furthermore, if we assume that knowledge is closed under multi-premise closure I should be able to know the conjunctive claim that tickets #1 & #2 &...#n-1 have lost, given that ticket #n won, in which case I would be able to deduce that ticket #n won, but clearly one cannot know which ticket won simply on the basis of the odds involved![13]

Even though the proposition that [tickets #1 & #2 &...#n-1 lost] is false in most nearby possible worlds, since in most possible worlds some other ticket than #n won, my belief that [tickets #1 & #2 &,...#n-1 lost] does satisfy SAFETY^S. After all, in all possible worlds where I *deduce* the conjunction from known premises, the whole conjunction is true, and hence there are no relevant nearby possible worlds where my belief is false. It could be objected that we are not infallible in our deductive capacities and that we sometimes make mistakes, and that therefore there is always a slight chance to infer something that is false from true premises when attempting to deduce a conclusion from a set of premises.[14] However, given that safety conditions are generally relativized to the method of belief-formation that the subject uses in the actual world, it is fair to demand that we keep constant the fact that the subject deduced the conclusion rather than that she merely inferred it.

The above paragraphs highlight some of the unwanted consequences of advancing the idea that we can know lottery propositions. If we grant that a subject can know of her own ticket that it

---

[13] According to multi-premise closure "[f]or all propositions P_1, … , P_n, Q, and all subjects s, if s knows each of P_1, … , P_n, and s comes to believe Q solely based on competent deduction from P_1, … , P_n, while retaining knowledge of each of P_1, … , P_n throughout, then s knows Q" (Lasonen-Aarnio 2008: 158). Lasonen-Aarnio argues that multi-premise closure and single-premise closure comes in a package: either we should keep both or reject both.
[14] Lasonen-Aarnio (2008) provides forceful arguments for the conclusion that weak safety conditions violate single-premise closure if the method of belief formation in paradigmatic cases of deductive knowledge is taken to be fallible.

has lost there will not be any principled reason why that subject should not be able to know the same of every ticket that has lost. If we assume multi-premise closure the subject can then deduce which ticket won the lottery. If we reject multi-premise closure, we should also reject single-premise closure, as Lasonen-Aarnio (2008) argues. If we claim that the safety condition should not be relativized to competent deduction, but rather to inference in cases of deductive knowledge, weak safety conditions such as SAFETY^S will violate single-premise closure (Lasonen-Aarnio 2008). Given that Sosa clearly endorses the idea that knowledge is closed under deductive inferences (1999; 2004: 304), he should accept both closure principles and claim that SAFETY^S should be relativized to competent deduction rather than mere inference. But if Sosa accepts multi-premise closure and that we can know that that a lottery ticket has lost simply on the basis of the odds involved, absurd consequences follow.

Given that there are versions of the safety condition that do not share these unwanted consequences, we have good reason to think that the kind of safety condition that is entailed by Sosa's virtue-theoretic conditions is not the safety condition that we have in mind when we require that knowledge has to be safe from error. For instance, a subject who believes that her lottery ticket has lost simply on the basis of the odds involved does not satisfy the safety condition that was given in section 2, since there is a very close nearby possible world where the subject continues to believe that she lost on the same basis and yet her belief is false. After all, the world where the subject wins the lottery is very similar to the actual world since only a few coloured balls would have to have dropped out of the lottery machine in a slightly different configuration (Pritchard 2007: 292). Crucially, the formulation of the safety condition as it was given in section 2 is motivated by the idea that knowledge is incompatible with epistemic luck

(Pritchard 2007). In order to successfully argue for the entailment thesis, the proponent of robust virtue epistemology must prove that the satisfaction of the ability condition entails the satisfaction of the right kind of safety condition. If the satisfaction of the ability condition does not entail the satisfaction of right formulation of the safety condition, then the motivation for the incompatibilist thesis diminishes.

4. CONCLUDING REMARKS

We have found that the arguments that have been laid down for the entailment thesis are ultimately unsatisfactory. They fail to show that the satisfaction of the ability condition entails the satisfaction of the right kind of safety condition, the one that we have in mind when we require our beliefs to be safe from error in order for those beliefs to amount to knowledge.[15]

REFERENCES

Ball, B. 2016. 'Knowledge, Safety, and Questions.' *Unisinos Journal of Philosophy,* 17: 58-62.
Bird, A. 2007. 'Justified Judging.' *Philosophy and Phenomenological Research,* 74: 81-110.
Bogardus, T. 2014. 'Knowledge Under Threat.' *Philosophy and Phenomenological Research,* 88: 289-313.
Carter, J. A. 2016. 'Robust Virtue Epistemology As Anti-Luck Epistemology: A New Solution.' *Pacific Philosophical Quarterly,* 97: 140-155.
Coffman, E. J. 2007. 'Thinking about luck.' *Synthese,* 158: 385-398.
Comesana, J. 2005. 'Unsafe knowledge.' *Synthese,* 146: 395-404.
Dretske, F. 1971. 'Conclusive Reasons.' *Australasian Journal of Philosophy,* 49: 1-22.
Gaultier, B. 2014. 'Achievements, Safety and Environmental Epistemic Luck.' *Dialectica,* 68: 477-497.
Greco, J. 2010. *Achieving Knowledge: A Virtue Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
Greco, J. 2012. 'A (Different) Virtue Epistemology.' *Philosophy and Phenomenological Research,* 85: 1-26.

Greco, J. 2016. 'Knowledge, Virtue, and Safety.' In M. A. Fernández Vargas (ed.), *Performance Epistemology - Foundations and Applications*, pp. 51-61. Oxford: Oxford University Press.

Hawthorne, J. and Lasonen-Aarnio, M. 2009. 'Knowledge and Objective Chance.' In P. Greenough & D. Pritchard (eds), *Williamson on Knowledge*, pp. 92-108. Oxford: Oxford University Press.

Hirvelä, J. 2017a. 'Global Safety: How to Deal with Necessary Truths.' *Synthese*.

Hirvelä, J. 2017b. 'Is it Safe to Disagree?' *Ratio,* 30: 305-321.

Ichikawa, J. 2014. 'Justification Is Potential Knowledge.' *Canadian Journal of Philosophy,* 44: 184-206.

Kelp, C. 2013. 'Knowledge: The Safe-Apt View.' *Australasian Journal of Philosophy,* 91: 265-278.

Lackey, J. 2007. 'Why We Don't Deserve Credit for Everything We Know.' *Synthese,* 158: 345-361.

Lackey, J. 2008. 'What Luck Is Not.' *Australasian Journal of Philosophy,* 86: 255-267.

Lackey, J. 2009. 'Knowledge and Credit.' *Philosophical Studies,* 142: 27-42.

Lasonen-Aarnio, M. 2008. 'Single Premise Deduction and Risk.' *Philosophical Studies,* 141: 157-173.

Littlejohn, C. 2014. 'Fake Barns and False Dilemmas.' *Episteme-a Journal of Individual and Social Epistemology,* 11: 369-389.

Luper, S. 2003. 'Indiscernibility Skepticism.' In S. Luper (ed.), *The Skeptics: Contemporary Essays*, pp. 183-202. Aldershot: Ashgate.

Miracchi, L. 2015. 'Competence to Know.' *Philosophical Studies,* 172: 29-56.

Neta, R. and Rohrbaugh, G. 2004. 'Luminosity and the Safety of Knowledge.' *Pacific Philosophical Quarterly, 85*: 396-406.

Nozick, R. 1981. *Philosophical Explanations*. Oxford: Oxford University Press.

Pritchard, D. 2005. *Epistemic Luck*. Oxford: Oxford University Press.

Pritchard, D. 2007. 'Anti-Luck Epistemology.' *Synthese,* 158: 277-297.

Pritchard, D. 2010. 'Knowledge and Understanding.' In D. Pritchard, A. Millar and A. Haddock (eds), *The Nature and Value of Knowledge*, pp. 5-88. Oxford: Oxford University Press.

Pritchard, D. 2012. 'Anti-Luck Virtue Epistemology.' *Journal of Philosophy,* 109: 247-279.

Riggs, W. 2007. 'Why Epistemologists Are so Down on Their Luck.' *Synthese,* 158: 329-344.

Sainsbury, R. 1997. 'Easy Possibilities.' *Philosophy and Phenomenological Research,* 57: 907-919.

Sosa, E. 1999. 'How to Defeat Opposition to Moore.' *Philosophical Perspectives,* 13: 141-154.

Sosa, E. 2004. 'Replies.' In J. Greco (ed.), *Ernest Sosa and his Critics*. Oxford: Blackwell.

Sosa, E. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge: Volume I*. Oxford: Oxford University Press.

Sosa, E. 2009. *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume II*. Oxford: Oxford University Press.

Sosa, E. 2015. *Judgment and Agency*. Oxford: Oxford University Press.

Turri, J. 2011. 'Manifest Failure: The Gettier Problem Solved.' *Philosophers Imprint,* 11: 1-11.

Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, T. 2009. 'Reply to John Hawthorne and Maria Lasonen-Aarnio.' In P. Greenough and D. Pritchard (eds), *Williamson on Knowledge*, pp. 313-329. Oxford: Oxford University Press.

Zagzebski, L. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press.

Zagzebski, L. 2003. 'The Search for the Source of the Epistemic Good'. *Metaphilosophy*, 34:12-28.